PROCEEDINGS OF SCIENCE

arXiv:1507.03908v1 [hep-ph] 14 Jul 2015

Fast evaluation of theoretical uncertainties with Sherpa and MCgrid

---

ZU-TH 21/15, MCnet-15-18

Enrico Bothmann[*]
II. Physikalisches Institut, Georg-August-Universität Göttingen, 37077 Göttingen, Germany
E-mail: enrico.bothmann@phys.uni-goettingen.de

Marek Schönherr
Institut für Theoretische Physik, Universität Zürich, 8057 Zürich, Switzerland
E-mail: marek.schoenherr@physik.uzh.ch

Steffen Schumann
II. Physikalisches Institut, Georg-August-Universität Göttingen, 37077 Göttingen, Germany
E-mail: steffen.schumann@phys.uni-goettingen.de

The determination of theoretical error estimates and PDF/α_s-fits require fast evaluations of differential cross sections for varied QCD input parameters. These include PDFs, the strong coupling constant α_S and the renormalisation and factorisation scales. Beyond leading order QCD, a full dedicated calculation for each set of parameters is often too time-consuming, certainly when performing PDF-fits. We report on two methods to overcome this issue for any QCD NLO calculation: The novel event-reweighting feature in SHERPA and the automated generation of interpolation grids using the recently introduced MCGRID interface. For MCGRID we present the newly added support for FASTNLO tables and highlight some future developments.

XXIII International Workshop on Deep-Inelastic Scattering
April 27–May 1, 2015
Dallas, Texas

---

[*] Speaker.

© Copyright owned by the author(s) under the terms of the Creative Commons Attribution 4.0 licence.

1. Introduction

In the last decade basically the complete automation of QCD NLO calculations has been achieved [1]. The predictions of these calculations depend on the QCD input parameters, such as the renormalisation scale μ_R, the factorisation scale μ_F, the strong coupling $\alpha_s(\mu_\mathrm{R}^2)$ and the parton density functions (PDFs) $f_a(x,\mu_\mathrm{F}^2)$ for the gluon and the (anti-)quarks.

These input parameters need to be varied when we want to calculate theory uncertainties, e.g. when we follow the convention to vary the scales with factors 0.5 and 2 to estimate the effect of higher-order corrections. But for QCD cross section calculations beyond the LO it usually takes days to achieve a precision that is comparable to data of current experiments, e.g. at the Large Hadron Collider. So a full re-evaluation for a given set of input parameter values uses up lots of CPU cycles. This problem is even more severe for another application of parameter variation: Fits of α_s and PDFs, where many re-evaluations are needed before the fit converges. Certainly for PDFs, which have many degrees of freedom in a fit, performing dedicated calculations for each PDF candidate value becomes impossible.

The solution is to identify how a prediction depends on the QCD input parameters. Then we can reuse the independent and CPU-intense parts for subsequent re-evaluations. This has been achieved for those types of QCD calculations, where the input parameters (or a known function of them) enter the calculation as prefactors to the identified independent parts [2].

There are two approaches to re-evaluate these independent parts for another choice of input parameter values. One is direct reweighting: We present here a novel feature of the 2.2.0 release of the Monte-Carlo event generator SHERPA [3] for performing an on-the-fly reweighting. Another approach is to create an interpolation grid, which allows for even faster re-evaluations. Currently there are two implementations for those grids, APPLGRID [4] and FASTNLO [5], and two interfaces for their automated creation using NLO events, AMCFAST [6] and MCGRID [7]. Here we discuss the new MCGRID 2.0 release and present its newly added support for FASTNLO and for filling the fixed-order expansion of MC@NLO calculations. We also present a method for quantifying the dependency of a parton shower algorithm on the QCD input parameters. This is a first step towards the goal of correctly taking care of these dependencies in a reweighting or an interpolation grid creation.

2. QCD NLO Events

Consider the structure of an NLO QCD calculation using Catani-Seymour subtraction [8], which consists of the following building blocks: The Born term (B), the virtual correction plus I-operator of the integrated subtraction terms (VI), the K- and P-operators of the integrated subtraction terms including the collinear factorisation remainders (KP), and the real-emission (R) and correlated subtraction (D_S) terms:

$$\begin{aligned}\langle O\rangle^{\mathrm{NLO}} = \lim_{N\to\infty}\frac{1}{N_\mathrm{trial}}\Bigg\{&\sum_i^{N_B}\Big\{\mathrm{B}(\Phi_{B,i})+\mathrm{VI}(\Phi_{B,i})+\mathrm{KP}(\Phi_{B,i},x'_{a/b})\Big\}\,O(\Phi_{B,i})\\ &+\sum_i^{N_R}\Big[\mathrm{R}(\Phi_{R_i})\,O(\Phi_{R_i})-\sum_j \mathrm{D}_S(\Phi_{B,j,i}\cdot\Phi^j_{R_i|B})\,O(\Phi_{B,j,i})\Big]\Bigg\},\end{aligned} \tag{2.1}$$

with $N_{\text{trial}} = \sum_i^N n_{\text{trial}\,i}$ the sum of the number of trials for each event and with the observable cuts and projections $O(\Phi)$, where Φ gives the phase space point. The $O(\Phi)$ are only dependent on the final state and so are independent of the QCD input parameters. We will now make explicit the dependency structure wich respect to the input parameters of each of the building blocks B, VI, KP and D_S. The simplest are the "Born-like" terms $X = \mathrm{B}, \mathrm{R}, \mathrm{D}_S$, where a combination of QCD input parameters occurs as a total prefactor, as is well-known from the LO factorisation theorem:

$$X(\Phi) = f_a(x_a, \mu_F^2)\, f_b(x_b, \mu_F^2)\, \alpha_s^{n+p}(\mu_R^2)\, X'(\Phi)\,, \tag{2.2}$$

where n is the number of strong vertices in the LO diagrams and p is either 0 (B) or 1 (R, D_S). Here and in the following the primed quantity (X') on the right has no dependency on the QCD input parameters. We call it the weight of the corresponding term (X). As one can see, for the evaluation of the PDFs the longitudinal momentum fractions $x_{a/b}$ have to be known in addition to the weight. The structure of the VI term is more complicated due to the occurring renormalisation scale logarithms:

$$\mathrm{VI}(\Phi_B) = f_a(x_a, \mu_F^2)\, f_b(x_b, \mu_F^2) \left[\alpha_s^{n+1}(\mu_R)\, \mathrm{VI}'(\Phi_B) + \left(c_R^{(0)} l_R + \tfrac{1}{2} c_R^{(1)} l_R^2 \right) \right], \tag{2.3}$$

with coefficients $c_R^{(i)}$ and scale logarithms $l_R = \log \frac{\mu_R^2}{\tilde{\mu}_R^2}$, where $\tilde{\mu}_R$ is the scale at which the $c_R^{(i)}$ had been evaluated in the first place. So for re-evaluating the VI term, not only the weight VI' and the momentum fractions $x_{a/b}$ need to be provided, but also the coefficients $c_R^{(i)}$ and the original scale $\tilde{\mu}_R$. The last missing building block are the KP terms:

$$\begin{aligned} \mathrm{KP}(\Phi_B, x'_{a/b}) = \alpha_s^{n+1}(\mu_R) \Big[& f_b(x_b, \mu_F^2) \left(f_a^q c_{F,a}^{(0)} + f_a^q(x'_a) c_{F,a}^{(1)} + f_a^g c_{F,a}^{(2)} + f_a^g(x'_a) c_{F,a}^{(3)} \right) \\ & + f_a(x_a, \mu_F^2) \left(f_b^q c_{F,b}^{(0)} + f_b^q(x'_b) c_{F,b}^{(1)} + f_b^g c_{F,b}^{(2)} + f_b^g(x'_b) c_{F,b}^{(3)} \right) \Big]\,, \end{aligned} \tag{2.4}$$

with $x'_{a/b}$ being the ratio of $x_{a/b}$ before and after the initial-state branching, and the coefficients

$$c_{F,a/b}^{(i)} = c_{F,a/b}^{\prime(i)} + c_{F,a/b}^{\prime\prime(i)} l_F \quad \text{for} \quad i = 0, \ldots, 3 \quad \text{with} \quad l_F = \log \frac{\mu_F^2}{\tilde{\mu}_F^2}\,. \tag{2.5}$$

The c'_F and c''_F are independent of the QCD input parameters. Again, $\tilde{\mu}_F$ is the scale at which the c_F had been evaluated. The functions $f^{q/g}$ give the correct combination of PDFs for the possible initial-state branchings, depending on whether a quark or a gluon is split, cf. [7].

All of these weights, coefficients and kinematics information can be passed in a standard `HepMC` [9] event record: The record comes with a `HepMC::WeightContainer` object, which can be used to store all weights and coefficients. This must of course be supported by the event generator, as done by e.g. SHERPA [3].

With this information provided, the weights and coefficients can be combined with any arbitrary choice of QCD input parameters. The same is true for the NLO expansion of MC@NLO and (partly) for multi-leg merging calculations, but their structure will not be discussed here to keep this note short. It is not true however for parton showers, their inherent all-orders resummation structure renders such a direct decomposition to prefactors composed of input parameters and a fixed number of weights/coefficients impossible. This lack of a correct treatment for variations in the shower results in a caveat for PDF variations of multi-leg merging calculations: They are not fully exact, because for them the PDFs enter via the shower branching probabilities determining the rejection weights for emissions already at the matrix element level.

3. Sherpa-Internal Reweighting

The newly added internal reweighting feature of SHERPA 2.2.0 uses the extraction of weights and coefficients sketched in the previous section to perform an on-the-fly reweighting for variations of the QCD input parameters. Any number of combinations of scales and PDFs can be specified prior to the run. The results are inserted into the `HepMC::WeightContainer` object and thus can be either saved to disk, programmatically accessed or directly passed via the internal interface to the RIVET analysis framework [10]. LO, NLO, the fixed-order expansion of MC@NLO and/or multi-jet merged runs are supported. The parton shower can be enabled in the run, but its dependency on the QCD input parameters can not be resolved yet as discussed in the previous section. The reweighting is then incomplete and the resulting prediction is not exact. Figure 1 illustrates the validation for the internal reweighting approach. Shown is the PDF variation plot for the transverse-momentum distribution in Higgs-boson production at the LHC, comparing via a single run with reweighting, with dedicated event generation runs for the modified PDF sets. For this example, the reweighting run takes approximately 4 times longer than a usual run, but generates the same information (with less statistical fluctuation) as 52 dedicated runs. In this example (and all the following ones) the input PDF has been interfaced using LHAPDF 6 [11].

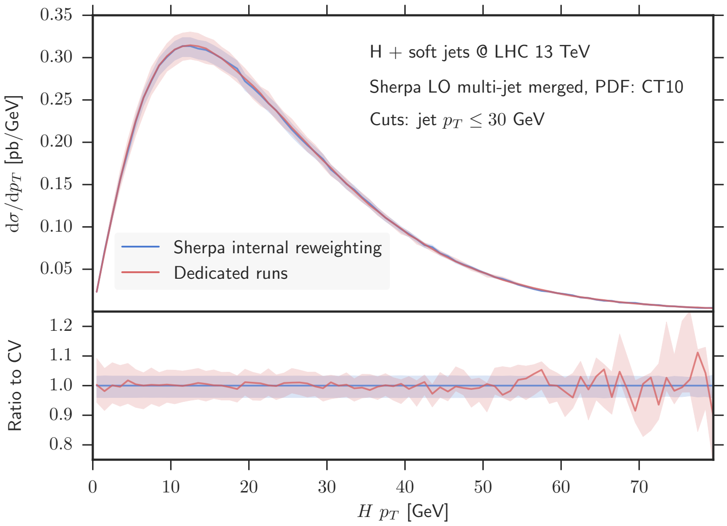

Figure 1: Validation for the new internal reweighting feature of SHERPA 2.2.0. A PDF band for the CT10 set [12] is generated using a single run with reweighting enabled to generate all PDF error eigenvalues at the same time. This is compared with the band for the same PDF set, but now doing a dedicated event generation for each set, i.e. without any reweighting.

4. Automated Grid Creation using MCgrid

The internal reweighting feature of SHERPA is useful for evaluating theory uncertainties for an observable. But when speed becomes very critical, because the number of re-evaluations needed becomes large (e.g. when fitting PDFs), more specialised methods must be used to reduce CPU time drastically. One established approach is based on cross section interpolation grids that represent a given differential cross section binned in the incoming partonic momentum fractions $x_{a/b}$ and the process scale $\mu_R^2 = \mu_F^2 = Q^2$. There are two implementations for those grids called APPLGRID [4] and FASTNLO [5]. While interpolating grids are defined only for a specific final state observable, they allow for a very quick re-evaluation of the respective cross section and have considerably reduced disk space requirements, compared to full event data. Until recently, interfaces to fill these grids were only available for process-specialised code. This has changed, as two interfaces to multi-purpose Monte Carlo event generators are now available.

One is called AMCFAST [6]: It interfaces LO and NLO predictions using the FKS scheme [13] generated with MADGRAPH5_aMC@NLO [14] to APPLGRID.

The second one is called MCGRID, which is implemented as a plugin to RIVET [10], such that (new or existing) RIVET analyses can be modified to create and fill interpolations grids. MCGRID reads in the events in the `HepMC` [9] event data format, such that in principle every Monte Carlo event generator can be used to produce grids via MCGRID. However, all required event weights must be inserted into the weight vector of the `HepMC` event. Currently, only the format that SHERPA writes into the weight vector is supported by MCGRID, which is based on the Catani-Seymour dipole subtraction scheme as illustrated in sec. 2. As for AMCFAST, since the first release of MCGRID [7], APPLGRID interpolation grids can be filled with LO and NLO events. Here we report on two newly added features in the 2.0 release of MCGRID: the support of FASTNLO interpolation grids and the exact treatment of the fixed-order expansion of MC@NLO events. This is the first time the FASTNLO package can be interfaced to a multi-purpose Monte Carlo event generator. With APPLGRID and FASTNLO, all currently available interpolation tools for fixed-order QCD cross sections can now be used in conjunction with MCgrid. For more details on the revisions made in MCGRID 2.0, see [15]. Figure 2 illustrates for the case of inclusive jet production at LHC energies, that FASTNLO and APPLGRID interpolation grids created with MCGRID both reproduce the distribution of the observable with per-mille level accuracy, much smaller than experimental and inherent theoretical errors.

5. Application: Residual PDF Dependencies of the Parton Shower

As has been explored earlier [16], varying the QCD input parameters in parton showers can have a sizeable effect on the result. The possibility to correctly vary the full fixed-order expansion of an MC@NLO calculation—via the internal reweighting of SHERPA or using interpolation grids created with MCGRID—allows for the quantification of the residual dependencies of the parton showers that are beyond the fixed-order approximation. Focusing on the PDF being used, the method would go as follows: First, generate PDF error bands for a given observable for two different PDF sets A and B. Use the event generation for the central value member of A to obtain distributions for all members of B via internal reweighting or an interpolation grid. Now the

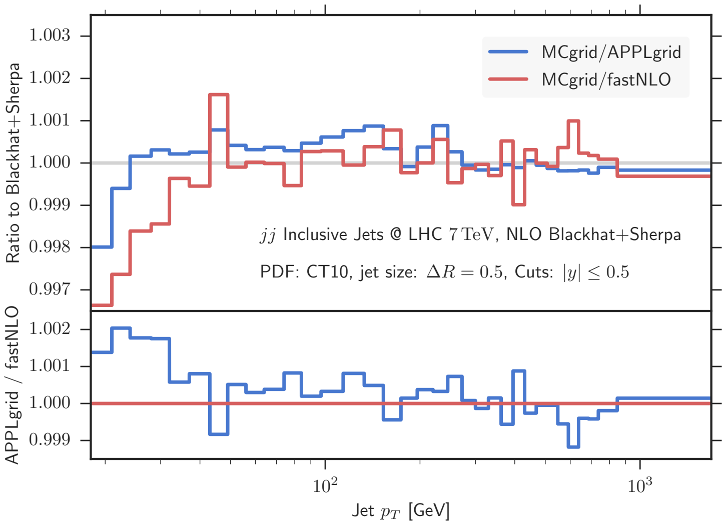

Figure 2: Validation for newly added support for FASTNLO grids in the 2.0 release of MCGRID. An inclusive jet p_T distribution is generated using a single run. At the same time, interpolation grids for FASTNLO and APPLGRID are written out via MCGRID. Their convolution with the same QCD input parameters used in the run are compared with the results from the direct run. The deviation is of a per-mille order. This has been achieved without any optimisations of the grid architecture (number of nodes/interpolation order).

reweighted/interpolated PDF band can be compared to the directly generated band for B. Because the PDF has been varied everywhere but in the parton shower, any (statistically significant) deviation in this comparison between the two B bands would indicate the associated observable bin's sensitivity to the PDF input of the parton shower. Finding sizeable deviations can motivate an implementation for tracking the dependencies of the parton shower. The sensitive bins would then serve as test cases to check if such an implementation works as expected, as the deviations must vanish then (or shrink if an approximation is used as e.g. only tracking the first parton shower emission). Interpolation grids filled with additional information for the parton shower would help in the fitting of PDFs to data that are not appropriately described by fixed-order calculations.

In fig. 3 we exemplify such a study with W boson production at the LHC, considering the p_T distribution of the boson. The grid has been filled with residual dependencies on the PDF set that has been used during its creation: the CT10 [12] central value member. When convoluting it with the MSTW pdf set [17] and comparing it with a set of dedicated runs, we observe deviations in the low $p_T \lesssim 5\,\text{GeV}$ region. This corresponds to a shift of the Sudakov peak position. The peak's shape and position is sensitive to multiple emissions and to the behaviour of the PDFs at very small scales. Therefore a careful consideration of these sensitivities must precede a definite conclusion.

6. Conclusion

We have presented two novel methods for fast re-evaluations of fixed-order cross sections suitable for the determination of theoretical uncertainties or PDF/α_s-fits in an automated way. Although both methods are based on reusing all the parts of the calculation that are not dependent on QCD input parameters, they do not require event data to be written out.

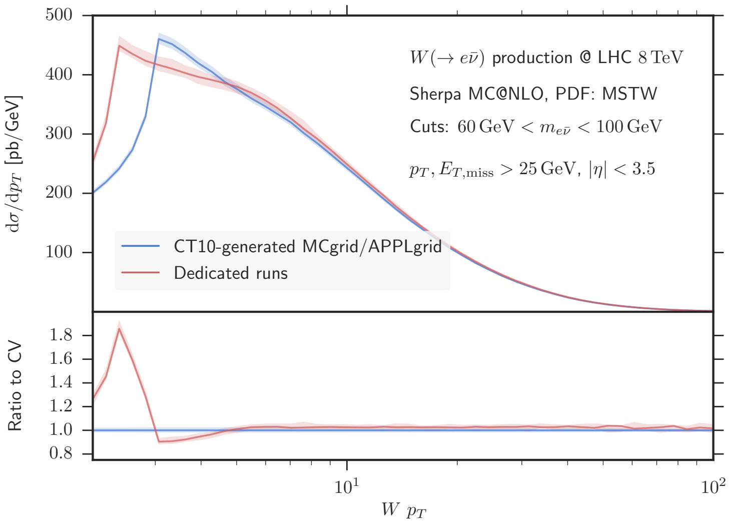

Figure 3: A Study of the residual PDF dependence of the parton shower of a grid that has been produced in an MC@NLO event generation with the central set of CT10 PDF. The MSTW-convolution of the grid deviates from the dedicated MSTW event generation in the low $p_T \lesssim 5\,\text{GeV}$ region.

The internal reweighting feature of SHERPA 2.2.0 is doing the reweighting on-the-fly, i.e. on an event-per-event basis. This is especially useful for simultaneously generating all distributions needed for specifying the theory errors for a given observable. MCGRID on the other hand is an interface for automated grid production using general-purpose event generators and RIVET. Because interpolation grids cut down the CPU time needed for a re-evaluation to the order of milliseconds, they are apt for PDF fitting.

The support of both methods for filling the fixed-order expansion of an MC@NLO calculation gives rise to a way to study the dependencies of the parton shower on the QCD input parameters. We have exemplified this for the transverse momentum distribution of the *W* boson.

In the future we plan to extend these studies, aiming to eventually implementing a method for taking into account (approximate) parton shower dependencies for the internal reweighting of SHERPA and for the creation of interpolation grids via MCGRID. Other features we want to add to MCGRID are NNLO predictions and the support for multi-jet merging calculations.

Acknowledgments

We acknowledge financial support from BMBF under contract 05H12MG5, the Swiss National Foundation (SNF) under contract PP00P2–128552 and from the EU MCnetITN research network funded under Framework Programme 7 contract PITN-GA-2012-315877.

References

[1] T. Carli, K. Rabbertz and S. Schumann, arXiv:1506.03239 [hep-ex].

[2] Z. Bern et al., Comput. Phys. Commun. **185** (2014) 1443.

[3] T. Gleisberg et al., JHEP **0902** (2009) 007.

[4] T. Carli et a., Eur. Phys. J. C **66**, 503 (2010).

[5] T. Kluge, K. Rabbertz and M. Wobisch, hep-ph/0609285.

[6] V. Bertone, R. Frederix, S. Frixione, J. Rojo and M. Sutton, JHEP **1408** (2014) 166.

[7] L. Del Debbio, N. P. Hartland and S. Schumann, Comput. Phys. Commun. **185** (2014) 2115.

[8] S. Catani and M. H. Seymour, Nucl. Phys. B **485** (1997) 291 [Nucl. Phys. B **510** (1998) 503].

[9] M. Dobbs and J. B. Hansen, Comput. Phys. Commun. **134** (2001) 41.

[10] A. Buckley et al., Comput. Phys. Commun. **184** (2013) 2803.

[11] A. Buckley et al., Eur. Phys. J. C **75** (2015) 3, 132.

[12] H. L. Lai et al., Phys. Rev. D **82** (2010) 074024.

[13] R. Frederix et al., JHEP **0910** (2009) 003.

[14] J. Alwall et al., JHEP **1106** (2011) 128.

[15] E. Bothmann, N. P. Hartland and S. Schumann, 10.1016/j.cpc.2015.07.002.

[16] S. Gieseke, JHEP **0501** (2005) 058.

[17] A. D. Martin et al., Eur. Phys. J. C **63** (2009) 189.